Immediate Observation in Mediated Population Protocols

Tobias Prehn

Modelle und Theorie Verteilter Systeme
TU Berlin, Germany

`tobias.prehn@tu-berlin.de`

Myron Rotter

TU Berlin, Germany

`m.rotter@campus.tu-berlin.de`

In this paper we analyze the computational power of variants of population protocols (PP), a formalism for distributed systems with anonymous agents having very limited capabilities. The capabilities of agents are enhanced in mediated population protocols (MPP) by recording the states in the edges of the interaction graph. Restricting the interactions to the communication model of immediate observation (IO) reduces the computational power of the resulting formalism. We show that this enhancement and restriction, when combined, yield a model (IOMPP) at least as powerful as the basic PP. The proof requires a novel notion of configurations in the MPP model allowing differentiation of agents and uses techniques similar to methods of analyzing encoding criteria, namely operational correspondence. The constructional part of the proof is generic in a way that all protocols can be translated into the new model without losing the desirable properties they might have besides a stable output. Furthermore, we illustrate how this approach could be utilized to prove our conjecture of IOMPP model being even as expressive as the MPP model. If our conjecture holds, this would result in a sharp characterization of the computational power and reveal the nonnecessity of two-way communication in the context of mediated population protocols.

1 Introduction

Population protocols have been introduced in 2004 as a computational model for passively mobile finite state sensors by Angluin et al. [2, 3]. They feature a finite state space, making them suitable for computation units with very limited capabilities and full anonymity, resulting directly from this restriction. Since the number of possible states that each agent could be in may not grow with the number of participating agents, there is no space for memorizing the ids of already met communication partners or similar constructs. Therefore, the outcome of any binary communication does not depend on whether the participants have communicated before. Another feature is the fully distributed approach of the base version for population protocols that does not need a base station, leader, or scheduler of any kind. The impact of such extensions has been studied [5, 7, 1].

It is well known that predicates computable by population protocols are exactly the semilinear predicates. The first study on the computational power of this model was in 2007 by Angluin et al. [4]. In this context, also several different communication patterns have been modeled in population protocols and their computational power have been studied as well. One of those mechanisms has been the immediate observation model, which is a special kind of one-way communication as opposed to the two-way communication that comes with the base model. The idea is that an agent may observe another agent without it noticing being observed. Clearly the observed agent cannot change its state in such an interaction whereas the observer can use the information given by its own and the observed agent's state. In contrast to stronger mechanisms no synchronization between the communication partners is needed. Consequently the communication in such a model is asynchronous and applicable to a broader variety of systems. With the fully distributed setting in mind the immediate observation communication seems to be a desirable feature. But these qualities come with a price. Protocols with this limitation to the

communication can only compute predicates in COUNT$_*$, i.e. predicates that count multiplicities of input values and compare them to previously given thresholds.

Another approach to altering population protocols has been the work of Michail et al. [12]. In their model agents are allowed to store distinct information for different communication partners. To achieve this they extended the base formalism by states for each pair of agents residing in the edges of the interaction graph. Some previous work on mediated population protocols modelled directed interaction graphs with one state per edge [12] and others used undirected graphs where each edge has a state for each of its two endpoints [8]. This extension is a reasonable compromise between maintaining the anonymity of each agent and being able to memorize the already met communication partners. An agent is capable of telling an agent, that it has not yet communicated with, apart from one it has already met. But two other agents being in the same state and with the same communication history are still indistinguishable. Aside from this the edge states can be used for storing several other information. This mediated population protocols are able to compute all symmetric predicates in $NSPACE(n^2)$.

We now present a model in this paper that combines the extended storage possibilities of mediated population protocols and the limited communication model of immediate observation protocols. The computational power of our resulting formalism has to be studied as it is unclear how this extension and restriction interact.

In section 2 the basic formalisms and existing models are defined. We are using a representation of population that allows the distinction of agents from a global point of view. This does not interfere with the anonymity of the agents and is for analysis purposes only. Based on this we define our model of immediate observation mediated population protocols (IOMPP) in section 3. Subsequently we study the computational power of our model in section 4. We take the approach simulating population protocols in immediate observation mediated population protocols in 3.1. Additionally, we give a translation of configurations from one model to the other and define criteria such a translation has to meet for it to express desirable attributes in 4.1. Our work is inspired by and makes use of the encodability criteria stated by Gorla in [11]. To the best of our knowledge this technique is novel to population protocols in the way we utilize it in 4.2 to prove that immediate observation mediated population protocols are at least as expressive as the base model of population protocols. In 4.3 we conjecture that our approach could also be used to show that immediate observation does not restrict the computational power of mediated population protocols. We conclude our paper by a discussion on the given results and possible application of the used techniques in section 5. We also give an outlook on future work and open questions.

2 Technical preliminaries

First we introduce populations, which form the base of all population protocols. They are often modeled as multisets to emphasize the indistinguishability of the participating agents. We will use vectors where each entry represents the state of a specific agent, because we want to efficiently compare populations in different models from a global viewpoint. Note that this will not give agents a distinct id they could make use from their local point of view. A different kind of vector representation can be found in [10] and is not to be confused with ours. They use vectors where each entry describes for an agent state the multiplicities of agents in that state. Their vector representation efficiently stores sets of agents with their states, making it easy to identify equivalent protocol states. In contrast to this, our representation can be used to compare sets of agents in population protocols with sets of agents in extended variants like mediated population protocols.

Definition 1 (Populations). Let A be a nonempty finite set and $n \in \mathbb{N}$. Then A^n describes the set of n-tuples over A also referred to as vectors of length n. A population over A, denoted by $POP(A)$, is the set of all vectors of arbitrary but finite length over A. If $v \in A^n$ we use $|v| = n$ to denote the length of v and $(v)_i$ to reference the i^{th} element of v (with $i \in \mathbb{N}^+_{\leq n}$).

Next we define how calculations in population protocols are modeled from a local point of view. Each agent has the same set of states and rules how to change states according to a communication partner's state. Additionally, functions to map input to an initial state and a state to some output are defined to be identical for each agent.

Definition 2 (Population Protocols [6]). A population protocol P is a 5-tuple $P = (Q, \Sigma, I, O, \delta)$ where Q is a finite set of agent states, Σ a finite input alphabet, $I : \Sigma \to Q$ describes the input function, $O : Q \to \{0, 1\}$ is the output function, and $\delta : Q^2 \to Q^2$ is referred to as transition function and describes all possible pairwise interactions. We also making use of a set representation of $\delta \subseteq Q^4$ whenever we write $t = (p, q) \to_\delta (p', q')$ referring to a specific transition $t \in \delta$.

To analyze protocols we need a global perspective. Here states are configurations holding state information for each agent. This kind of global view has been used to study the computational power of population protocols [4].

Definition 3 (Global Protocols [4]). Let $P = (Q, \Sigma, I, O, \delta)$ be a population protocol. The global protocol to P is a 5-tuple $G_P = (\mathcal{C}, \Sigma, \mathcal{I}, \mathcal{O}, \longrightarrow)$ where $\mathcal{C} = POP(Q)$ is the set of configurations, i.e., vectors of agent states Q, $\mathcal{I} : POP(\Sigma) \to \mathcal{C}$ maps input vectors to initial configurations, $\mathcal{O} : \mathcal{C} \to \{0, 1, \bot\}$ maps configurations to outputs, and $\longrightarrow : \mathcal{C} \to \mathcal{C}$ is the global transition function with $\longrightarrow^*$ being its reflexive and transitive closure. For $C, C' \in \mathcal{C}$ it holds that $C \longrightarrow C'$ iff there is a transition $t \in \delta$ and $i, j \in \mathbb{N}^+$ with $i \neq j$ such that $t = ((C)_i, (C)_j) \to_\delta ((C')_i, (C')_j)$ and $(C)_k = (C')_k$ for every $k \in \mathbb{N}^+ \setminus \{i, j\}$. We also write $C \xrightarrow{t_{i,j}} C'$ and call agent i the *initiator* of t and j the *responder* or (if the state of i is not changed by t) *observer*. The global input function takes use of I to get an agent state for each single value in its input and $\mathcal{O}$ aggregates the outputs of the agents according to O. It holds that $\mathcal{O}(C) = x \in \{0, 1\}$ iff $O((C)_i) = x$ for each $i \in \mathbb{N}^+_{\leq |C|}$ and $\mathcal{O}(C) = \bot$ in every other case. When the underlying protocol P is clear from the context we often omit the index of G_P and simply state that G is the global protocol to P.

Based on a global protocol we can describe what it means for a protocol to compute some predicate. For this we need to define executions and fairness.

Definition 4 (Computation). Let $G = (\mathcal{C}, \Sigma, \mathcal{I}, \mathcal{O}, \longrightarrow)$ be a global protocol. A configuration $C \in \mathcal{C}$ is output stable with output $x \in \{0, 1\}$ iff $\mathcal{O}(C') = x$ for each $C' \in \mathcal{C}$ with $C \longrightarrow^* C'$. We call a sequence of configurations $C_0, C_1, C_2, \cdots \in \mathcal{C}$ with $C_i \longrightarrow C_{i+1}$ for each $i \in \mathbb{N}$ an execution. An execution is fair iff for each $C \in \mathcal{C}$ with $C_i = C$ for infinitely many $i \in \mathbb{N}$ it holds that if there is a transition $C \longrightarrow C'$ then also $C_j = C'$ for infinitely many $j \in \mathbb{N}$. A population protocol P is well-specified if for each input *Inp*, it holds that all fair executions of P starting in $\mathcal{I}(Inp)$ reach a configuration that is output stable. P computes a predicate if this reached configuration is output stable with output 1 if *Inp* satisfies the predicate and with output 0 otherwise.

In the context of population protocols several communication mechanisms have been studied [4]. Immediate observation is one of those mechanisms. It reduces the class of computable predicates to predicates counting multiplicities of input values COUNT$_*$. To model this kind of communication, restrictions to the allowed form of transitions are made.

Definition 5 (Immediate Observation). Let $P = (Q, \Sigma, I, O, \delta)$ be a population protocol. P is an immediate observation protocol, if there is no transition that changes the state of the initiator. In other words, all transitions $t \in \delta$ have to be of the form $t = (p, q) \to_\delta (p, q')$.

Another extension to population protocols is the mediated variant. The idea is to introduce states in all edges of the communication graph. Since the most general graph is the complete graph, each pair of agents is given such a state. In the context of an immediate observation communication mechanism, it is not reasonable to assume a storage that both agents can write to. Therefore, we introduce a pair of edge states for each pair of agents. Edge states are always initialized with the same value.

Definition 6 (Mediated Population Protocols [12]). A mediated population Protocol P is a 7-tuple $P = (Q, \Sigma, S, s_0, I, O, \delta)$ where Q, Σ, I and O are analogous to population protocols. The set of edge states S includes the initial edge state $s_0 \in S$ and the transition function $\delta : (Q \times S)^2 \to (Q \times S)^2$ incorporates the edge states for each pair of agents.

Configurations in mediated population protocols cannot be represented by simple vectors. We need to introduce matrices as configurations containing the agent states on the diagonal and the states of the edge between agents a and b in fields $C_{a,b}$ (side of agent a) and $C_{b,a}$ (side of agent b).

Definition 7 (Mediated Populations). $A^{n \times n}$ describes the set of square matrices over A of size $n \times n$. A mediated population over A denoted by $POP_M(A)$ is the set of all matrices of arbitrary but finite length over A. If $m \in A^{n \times n}$ we use $(m)_{i,j}$ to reference the element of m at column i and row j with $i, j \in \mathbb{N}^+_{\leq n}$ and $|m| = n$ to denote the length as well as the height of a square matrix m.

We can now proceed with lifting our global protocol definitions to represent mediated population protocols as well.

Definition 8 (Global Protocols for Mediated Population Protocols). Let $P = (Q, \Sigma, S, s_0, I, O, \delta)$ be a mediated population protocol. The global protocol to P is again a 5-tuple $G = (\mathcal{C}, \Sigma, \mathcal{I}, \mathcal{O}, \longrightarrow)$. In contrast to global protocols for simple population protocols $\mathcal{C} = POP_M(Q)$ is the set of configurations and $\mathcal{I} : POP(\Sigma) \to \mathcal{C}$ maps input vectors to initial configurations, initializing the diagonal fields with the corresponding agent states and every other field with s_0. The output function $\mathcal{O} : \mathcal{C} \to \{0, 1, \bot\}$ ignores all fields not on the diagonal and $\longrightarrow : \mathcal{C} \to \mathcal{C}$ now also changes the respective edge states. For $C, C' \in \mathcal{C}$ it holds that $C \longrightarrow C'$ iff there is a transition $t \in \delta$ and $i, j \in \mathbb{N}^+$ with $i \neq j$ such that $t = ((C)_{i,i}, (C)_{i,j}, (C)_{j,j}, (C)_{j,i}) \to_\delta ((C')_{i,i}, (C')_{i,j}, (C')_{j,j}, (C')_{j,i})$ and $(C')_{k,l} = (C')_{k,l}$ for every $k, l \in \mathbb{N}^+ \setminus \{i, j\}$.

3 Modelling immediate observation in mediated population protocols

From the technical preliminaries in section 2 we can easily combine the models for mediated population protocols and immediate observation conform communication. We get our model of population protocols with two edge states in every edge, one per communication partner, and transitions that keeps the states of the initiator unaltered and changes the states of the observer.

Definition 9 (Immediate Observation Mediated Population Protocols). An immediate observation mediated population protocol P is a 7-tuple $P = (Q, \Sigma, S, s_0, I, O, \delta)$ where Q is a finite set of agent states, Σ a finite input alphabet, $I : \Sigma \to Q$ describes the input function, $O : Q \to \{0, 1\}$ is the output function, and $\delta : (Q \times S)^2 \to (Q \times S)^2$ is referred to as transition function and describes all possible pairwise interactions. We also making use of a set representation of $\delta \subseteq (Q \times S)^4$ whenever we write $t = (p, s, q, r) \to_\delta (p', s', q', r')$ referring to a specific transition $t \in \delta$. Since our model uses the immediate observation communication mechanism, all transitions $t = (p, s, q, r) \to_\delta (p', s', q', r')$ have to satisfy $p = p'$ and $s = s'$.

3.1 Simulating population protocols by immediate observation mediated population protocols

We can now simulate protocols in the basic population protocol model by immediate observation mediated population protocols. The main idea is to split every two-way communication with an initiator and a responder into 4 steps. Two steps are required to signal the request and the acknowledgement of a communication and another two steps are needed to finish the communication resolving all pending state changes. Additionally, a reset transition is given for the case of an unsuccessful communication.

Simulation 10. Let $P = (Q, \Sigma, I, O, \delta)$ be a population protocol. The following immediate observation mediated population protocol P' simulates the protocol P and is given by the tuple $(Q', \Sigma', S', s'_0, I', O', \delta')$ where

$$
\begin{aligned}
\Sigma' &:= \Sigma, \\
Q' &:= \{L, U\} \times Q, \\
S' &:= \{s_{init}, s_{ponr}\} \cup Q, \\
s'_0 &:= s_{init}, \\
I'(\sigma) &:= (U, I(\sigma)) \text{ for all } \sigma \in \Sigma, \\
O'(l, q) &:= O(q) \text{ for all } (l, q) \in Q'.
\end{aligned}
$$

The states of the agents are of the form (l, q) where $l \in \{L, U\}$ indicates whether the agent is locked or unlocked and $q \in Q$ is the computation state according to the original population protocol P. For easier referencing we call the first component of an agent state (l, q) the locking state l of this agent and the second component its computation state q. We use the formulation of an agent being locked whenever its locking state is L and say this agent is unlocked otherwise. W.l.o.g. we assume that $\{s_{init}, s_{ponr}\} \cap Q = \emptyset$. The input function I maps each input symbol $\sigma \in \Sigma$ to the state $(U, I(\sigma))$. The output function O' maps each state $(l, q) \in Q$ to $O(q)$ independent of the locking indicator l. We specify for each transition $t = (p, q) \to_\delta (p', q')$ of δ with $p, q, p', q' \in Q$ the following transitions for δ'.

$$
\begin{aligned}
t^{(1)} &= ((U, p), s_{init}, (U, q), s_{init}) &\to_{\delta'}& \quad ((U, p), s_{init}, (L, q'), q) & (1) \\
t^{(2)} &= ((L, q'), q, (U, p), s_{init}) &\to_{\delta'}& \quad ((L, q'), q, (L, p'), s_{ponr}) & (2) \\
t^{(3)} &= ((L, p'), s_{ponr}, (L, q'), q) &\to_{\delta'}& \quad ((L, p'), s_{ponr}, (U, q'), s_{init}) & (3) \\
t^{(4)} &= ((x, y), s_{init}, (L, p'), s_{ponr}) &\to_{\delta'}& \quad ((x, y), s_{init}, (U, p'), s_{init}) & (4) \\
& & & \quad \text{for every } (x, y) \in Q' \\
t^{(5)} &= ((x, y), z, (L, q'), q) &\to_{\delta'}& \quad ((x, y), z, (U, q), s_{init}) & (5) \\
& & & \quad \text{for every } (x, y) \in Q' \\
& & & \quad \text{and } z \in S' \setminus \{s_{ponr}\}
\end{aligned}
$$

The locking state of each agent prohibits simultaneous participation in several different communications. Whenever an agent took part in a two-way communication t in the original protocol, it could be the observer of a transition of type (1) in the simulation. A $t^{(1)}$ transition locks the observing agent and puts its old computation state in the edge state this agent controls on the edge with the observed agent. This has two reasons: First it signals the interest in a communication with the other agent and second it backups the old state for a potential future reset. If the other agent observes the change in the edge state, it may signal the acknowledgement of requested communication by locking itself, changing

its computation state according to the transition t and putting s_{ponr} in its edge state. This is achieved by transition $t^{(2)}$. Now the two agents have to reset their edge states to s_{init} and unlock themselves. The agent mimicking the responder of the original transition t starts by taking $t^{(3)}$, followed by the simulator of the original initiator taking $t^{(4)}$. If a communication was not successful, either because the initiators surrogate has taken an other transition with another agent in the meantime or because responder simulating agent observes its partner before it could acknowledge the communication, $t^{(5)}$ is taken. This transition assures that in the described cases an agent can give up on a communication attempt and reset its state, readying itself for another attempt, potentially with a different partner.

Note that if a transition $t = (p,q) \rightarrow_\delta (p,q')$ is already immediate observation compliant we do not need to add the whole set of transitions. We could instead add a slightly altered version of the original transition as follows.

$$t^{(6)} = ((U,p),s_{init},(U,q),s_{init}) \quad \rightarrow_{\delta'} \quad ((U,p),s_{init},(U,q'),s_{init}) \tag{6}$$

Clearly the result would be the same. This kind of transition is only needed if the simulation needs to be more efficient in the sense of steps needed to get to an output stable configuration. We will therefore omit this type of transitions in our analyses.

Observation 11 (Output Changing Transitions)**.** By definition of O' the output of an agent only depends on its computation state. As transitions $t^{(3)}$, $t^{(4)}$ do not change the computation states, only transitions $t^{(1)}$, $t^{(2)}$, and $t^{(5)}$ can have an impact on the output of an agent. Since Simulation 10 is an immediate observation protocol, this agent has to be the observer of such transitions.

Observation 12 (Number of Started Conversations)**.** Every agent has at most one started and not yet concluded conversation at any point in time. Starting a conversation by taking transition $t^{(1)}$ as responder brings an agent to a locked state. Therefore, no other conversation can be started or acknowledged by this agent until the conversation is concluded with transition $t^{(3)}$ or aborted with transition $t^{(5)}$. Acknowledging a conversation by taking transition $t^{(2)}$ as responder also brings an agent to a locked state. Again no other conversation can be started or acknowledged by this agent until the conversation is concluded with transition $t^{(4)}$ or aborted with transition $t^{(5)}$.

Observation 13 (Point of no Return)**.** Every occurrence of transition $t^{(2)}_{i,j}$ with acting agents i and j is eventually followed by transitions $t^{(3)}_{j,i}$ and $t^{(4)}_{i,j}$. After execution of $t^{(2)}_{i,j}$ agent i is locked with $q \in Q$ in its edge state to j and agent j is locked with s_{ponr} in its edge state to i. From Observation 12 we know that neither i nor j can be observer of any transition with some agent different from i and j. From the transitions with i and j only $t^{(3)}_{j,i}$ is enabled and will be taken at some point because of the fairness assumptions in population protocols. After that agent j is still locked with s_{ponr} in its edge state to i. Therefore, j can only be observer of transitions $t^{(3)}, t^{(4)}$, or $t^{(5)}$. Again from Observation 12 we know that only $t^{(4)}_{i,j}$ is possible.

4   Computational power

We now show that our model can compute all predicates computable in population protocols by giving a translation, that relates configurations from a protocol to configurations from its simulation representing the same state of computation. Additionally, we identify requirements imposed on such a translation to be helpful in proving the equality of computed predicates. We will ultimately show how this proof is executed.

4.1   Translation and criteria

We provide a translation, that constructs configurations in mediated population protocols from configurations of population protocols by giving all agents an unlocked state and setting all edges to the neutral s_{init} state.

Definition 14 (Translation of Configurations)**.** Let $P = (Q, \Sigma, I, O, \delta)$ be a population protocol and $P' = (Q', \Sigma', S', s_0', I', O', \delta')$ a immediate observation mediated population protocol constructed from P using Simulation 10. By $[\![\cdot]\!] : POP(Q) \to POP_M(Q')$ we denote the translation of configurations $C \in POP(Q)$ in the population protocol into configurations $D \in POP_M(Q')$ from the mediated population protocol. This translation is defined as follows:

$$
[\![(q_0, q_1, \ldots, q_n)]\!] =
\begin{pmatrix}
(U, q_1) & s_{init} & \cdots & s_{init} \\
s_{init} & (U, q_2) & \ddots & \vdots \\
\vdots & \ddots & \ddots & s_{init} \\
s_{init} & \cdots & s_{init} & (U, q_n)
\end{pmatrix}
$$

From the criteria for *good encodings* defined by Gorla [11] we adopt the notion of operational correspondence.

Definition 15 (Operational Correspondence)**.** A translation $[\![\cdot]\!] : POP(Q) \to POP_M(Q')$ is operationally corresponding if it is

(1) (operationally) complete, i.e., for all $C, C' \in POP(Q)$ with $C \longrightarrow^* C'$ it holds that $[\![C]\!] \longrightarrow^* [\![C']\!]$, and

(2) (operationally) sound, i.e., for all $C \in POP(Q)$ and $D \in POP_M(Q')$ with $[\![C]\!] \longrightarrow^* D$ there exists a $C' \in POP(Q)$ with $D \longrightarrow^* [\![C']\!]$ and $C \longrightarrow^* C'$.

If our translation in Def. 14 instantiated with concrete population protocol P and mediated population protocol P' is operationally corresponding, we get that every configuration reachable in P is also reachable in P' and vice versa.

Definition 16 (Input/Output Correspondence)**.** Let P be a population protocol, P' be a mediated population protocol and G, G' be the global protocols to P and P' respectively. A translation $[\![\cdot]\!] : POP(Q) \to POP_M(Q')$ is I/O corresponding if it is

(1) input corresponding, i.e., for all $V \in POP(\Sigma)$ it holds that $[\![\mathcal{I}(V)]\!] = \mathcal{I}'(V)$, and

(2) output corresponding, i.e., for all $C \in POP(Q)$ it holds that $\mathcal{O}(C) = \mathcal{O}'([\![C]\!])$.

Input/Output Correspondence gives us the assurance that input and output functions of the protocols related by a translation behave in a similar way. If it holds, translating an input configuration of the original protocol or directly using the input function of the corresponding mediated population protocol yields the same result. Additionally, configurations are always translated into configurations with the same output.

Definition 17 (Output Stability Preservation)**.** Let P be a population protocol, P' be a mediated population protocol and G, G' be the global protocols to P and P' respectively. A translation $[\![\cdot]\!] : POP(Q) \to POP_M(Q')$ is output stability preserving if for each $C \in POP(Q)$ it holds that $[\![C]\!]$ is output stable iff C is output stable.

From the output stability preservation we get that each output stable configuration is translated into a configuration also being output stable.

Lemma 18. *Let* $P = (Q, \Sigma, I, O, \delta)$ *be a population protocol and* $P' = (Q', \Sigma', S', s'_0, I', O', \delta')$ *a mediated population protocol. If there is a translation* $[\![\cdot]\!] : POP(Q) \to POP_M(Q')$ *that is operationally corresponding, input/output corresponding, and output stability preserving, then* P *and* P' *compute the same predicate.*

Proof. Since the translation is input corresponding every initial configuration in P has a corresponding initial configuration in P'. From the operational correspondence we know that a configuration is reachable from an initial configuration in P' iff it has a corresponding configuration reachable from the corresponding initial configuration in P. With output correspondence both configurations clearly have the same output and, since the translation is output stability preserving, the output is either stable in both configurations or in none. Therefore, the protocol P calculates the same semilinear predicate as P'. $\quad\square$

4.2 All semilinear predicates can be computed by immediate observation mediated population protocols

Lemma 19. *The translation* $[\![\cdot]\!]$ *given in Def. 14 is operationally corresponding for any population protocol* $P = (Q, \Sigma, I, O, \delta)$ *and the mediated population protocol* $P' = (Q', \Sigma', S', s'_0, I', O', \delta')$ *constructed using Simulation 10.*

Proof. To prove operational completeness assume that $C, C' \in POP(Q)$ with $C \longrightarrow^* C'$. Since $\longrightarrow^*$ is defined as reflexive-transitive closure of $\longrightarrow$ we get that $C_1, C_2, \ldots, C_n \in POP(Q)$ exist with $C \longrightarrow C_1 \longrightarrow C_2 \longrightarrow \ldots \longrightarrow C_n \longrightarrow C'$. We can always simulate a step $C_i \longrightarrow C_{i+1}$ in P by making 4 steps in P'. Assume that the step is due to transition $t \in \delta$ and agents at a and b are acting as initiator and responder respectively in $C_i \xrightarrow{t_{a,b}} C_{i+1}$. This can be simulated as $[\![C_i]\!] \xrightarrow{t^{(1)}_{a,b}} C_i^2 \xrightarrow{t^{(2)}_{b,a}} C_i^3 \xrightarrow{t^{(3)}_{a,b}} C_i^4 \xrightarrow{t^{(4)}_{b,a}} [\![C_{i+1}]\!]$ and thus $[\![C_i]\!] \longrightarrow^4 [\![C_{i+1}]\!]$ holds. Thus, $[\![C]\!] \longrightarrow^4 [\![C_1]\!] \longrightarrow^4 \ldots \longrightarrow^4 [\![C_n]\!] \longrightarrow^4 [\![C']\!]$ exemplifies $[\![C]\!] \longrightarrow^* [\![C']\!]$.

To prove operational soundness assume that $C \in POP(Q)$ and $D \in POP_M(Q')$ with $[\![C]\!] \longrightarrow^* D$. We construct C' as follows. For this assume $(D)_{i,i} = (l_i, d_i)$.

$$(C')_i = \begin{cases} g & \text{, if there is exactly one } j \in \mathbb{N}^+_{\leq |D|, \neq i} \text{ such that } (D)_{i,j} = g \text{ and } (D)_{j,i} \neq s_{ponr} \\ d_i & \text{, otherwise} \end{cases}$$

We can show that $D \longrightarrow^* [\![C']\!]$ by taking the appropriate transitions for all $i, j \in \mathbb{N}^+_{\leq |D|}$ with $i \neq j$ as follows.

$$\text{if } (D)_{i,j} = q \quad \text{and} \quad (D)_{j,i} = s_{ponr} \quad \text{take transitions } t^{(3)}_{j,i}, t^{(4)}_{i,j}$$
$$\text{if } (D)_{i,j} = s_{init} \quad \text{and} \quad (D)_{j,i} = s_{ponr} \quad \text{take transition } t^{(4)}_{i,j}$$
$$\text{if } (D)_{i,j} = q \quad \text{and} \quad (D)_{j,i} \neq s_{ponr} \quad \text{take transition } t^{(5)}_{j,i}$$
$$\text{otherwise} \qquad\qquad\qquad\qquad \text{do nothing}$$

Since $[\![C]\!] \longrightarrow^* D$ by assumption and $D \longrightarrow^* [\![C']\!]$ we get that $[\![C]\!] \longrightarrow^* [\![C']\!]$. We can construct $C \longrightarrow C_1 \longrightarrow C_2 \longrightarrow \ldots \longrightarrow C'$ from the path $[\![C]\!] \longrightarrow D_1 \longrightarrow D_2 \longrightarrow \ldots \longrightarrow [\![C']\!]$ as follows. Whenever there is a step $D_x \xrightarrow{t^{(2)}_{i,j}} D_y$ in this path, take transition $C_v \xrightarrow{t_{i,j}} C_w$ in the path of $C \longrightarrow^* C'$. By Observation 12 we get that each such step of type (2) is eventually followed by transitions of type (3) and (4) and since $[\![C']\!]$ has only edge states s_{init}, this has to happen before $[\![C']\!]$ is reached. Therefore, there exists a $C' \in POP(Q)$ with $D \longrightarrow^* [\![C']\!]$ and $C \longrightarrow^* C'$. $\quad\square$

Lemma 20. *The translation $[\![\cdot]\!]$ given in Def. 14 is I/O corresponding for any population protocol $P = (Q, \Sigma, I, O, \delta)$ and the mediated population protocol $P' = (Q', \Sigma', S', s'_0, I', O', \delta')$ constructed using Simulation 10.*

Proof. The input function I' of P' makes use of the input function I of P by putting an agent in the unlocked state $(U, I(\sigma))$ iff the same agent would be in state $I(\sigma)$ in P. Also every edge state is initialized with s_{init}. This matches $[\![\cdot]\!]$ that translates each state q to (U, q) and sets every edge state to s_{init}. Thus, $[\![\mathcal{I}(V)]\!] = \mathcal{I}'(V)$ and $[\![\cdot]\!]$ is input corresponding.

The output function O' of P' makes use of the output function O of P by ignoring the locking state of an agent and giving back the output of O for the computation state. As seen above, $[\![\cdot]\!]$ translates every agents state q to a tuple with q being the second component i.e. its computation state. For each agent a it holds that $O((C)_a) = O'\left((D)_{a,a}\right)$. As $\mathcal{O}$ aggregates the outputs of all agents, which are the same in P and P', $\mathcal{O}(C) = \mathcal{O}'([\![C]\!])$ and consequently $[\![\cdot]\!]$ is output corresponding. $\qquad\square$

Lemma 21. *The translation $[\![\cdot]\!]$ given in Def. 14 is output stability preserving for any population protocol $P = (Q, \Sigma, I, O, \delta)$ and the mediated population protocol $P' = (Q', \Sigma', S', s'_0, I', O', \delta')$ constructed using Simulation 10.*

Proof. Let C be a configuration that is output stable in P. Assume towards contradiction $[\![C]\!]$ is not output stable in P'. From the definition of output stability follows that a configuration D exists with $\mathcal{O}'([\![C]\!]) \neq \mathcal{O}'(D)$ and D is reachable from $[\![C]\!]$, i.e. there is a path $[\![C]\!] \longrightarrow D_1 \longrightarrow D_2 \longrightarrow \ldots \longrightarrow D$. W.l.o.g assume that D is the first such configuration in this path, i.e. $\mathcal{O}'([\![C]\!]) = \mathcal{O}'(D_i)$ for each such D_i. Since $\mathcal{O}'$ aggregates the values of O' for each agent and because P' is an immediate observation protocol, there has to be a single agent a that has changed its output because of the transition leading to D. From Observation 11 we know that this transition has to be of type (1), (2), or (5) and a has to be its observer.

If it is (1) or (2) we can construct a configuration in the same way as C' was constructed in the proof of operational soundness for Lemma 19. This C' is reachable from C in P and $[\![C']\!]$ is reachable from D in P'. Note that on the path from D to $[\![C']\!]$ agent a is never an observer of any transition with type (1), (2), or (5) and therefore does not change its output. Because our translation maintains outputs for each agent $O((C)_a) = O'\left(([\![C]\!])_{a,a}\right) \neq O'\left((D)_{a,a}\right) = O'\left(([\![C']\!])_{a,a}\right) = O((C')_a)$ holds. This results in $\mathcal{O}(C) \neq \mathcal{O}(C')$, a contradiction to C being output stable.

If the transition agent a took was of type (5), there has to be a configuration along the path from $[\![C]\!]$ to D where agent a took a transition of type (1). Observation 12 states that there has to be such a transition in advance and by the definition of $[\![\cdot]\!]$ this has to be after $[\![C]\!]$. But if the transition of type (5) changes the output of a, the corresponding type (1) transition must also have changed it, contradicting our assumption of D being the first configuration with an output different from $\mathcal{O}'([\![C]\!])$. This is due to a type (5) transition resetting the computation state of an agent back to the state it had before the corresponding type (1) transition.

For the other direction assume towards contradiction $[\![C]\!]$ is output stable in P' and C is not output stable in P. Then there exists a configuration $\overline{C}$ that is reachable from C with $\mathcal{O}(C) \neq \mathcal{O}(\overline{C})$. Because $[\![\cdot]\!]$ is operationally complete by Lemma 19 it holds that $[\![\overline{C}]\!]$ is reachable from $[\![C]\!]$. From the output correspondence of $[\![\cdot]\!]$ in Lemma 20 follows that $\mathcal{O}'([\![C]\!]) = \mathcal{O}(C) \neq \mathcal{O}(\overline{C}) = \mathcal{O}'([\![\overline{C}]\!])$. This is a contradiction to the output stability of $[\![C]\!]$ in P'. $\qquad\square$

Theorem 22. *For any population protocol $P = (Q, \Sigma, I, O, \delta)$, the immediate observation mediated population protocol P' constructed from P using Simulation 10 calculates the same semilinear predicate.*

Proof. We have shown that $[\![\cdot]\!]$ from Definition 14 is operationally corresponding, I/O corresponding and output stability preserving in Lemmas 19, 20, and 21. By Lemma 18 the statement directly follows. $\square$

Corollary 23. *Immediate observation mediated population protocols can compute every semilinear predicate and are therefore at least as expressive as population protocols.*

4.3 Immediate observation does not restrict the computational power of mediated population protocols

The approach in the previous sections can be used to prove that two-way communication in mediated population protocols does not add to the computational power of the model. To achieve this we define a simulation of mediated population protocols into the variant with immediate observation communication.

Simulation 24. Let $P = (Q, \Sigma, S, s_0, I, O, \delta)$ be a mediated population protocol. The following immediate observation mediated population protocol P' simulates the protocol P and is given by the tuple $(Q', \Sigma', S', s_0', I', O', \delta')$ where

$$
\begin{aligned}
\Sigma' &:= \Sigma, \\
Q' &:= \{L, U\} \times Q, \\
S' &:= (\{s_{init}, s_{ponr}\} \cup (Q \times S)) \times S, \\
s_0' &:= (s_{init}, s_0), \\
I'(\sigma) &:= (U, I(\sigma)) \text{ for all } \sigma \in \Sigma, \\
O'(l, q) &:= O(q) \text{ for all } (l, q) \in Q'.
\end{aligned}
$$

W.l.o.g. we assume that $\{s_{init}, s_{ponr}\} \cap (Q \cup S) = \emptyset$. In contrast to Simulation 10 the edge state has two components. The first component again signals the current state of the simulated communication and serves as a backup for the condition prior to the communication. Here we need to save both, computation state and edge state. The second component represents the actual edge state present in the original protocol. We specify for each transition $t = (p, r, q, s) \rightarrow_\delta (p', r', q', s')$ of δ with $p, q, p', q' \in Q$ and $r, s, r', s' \in S$ the following transitions for δ'.

$$
\begin{aligned}
t^{(1)} &= ((U, p), (s_{init}, r), (U, q), (s_{init}, s)) &\rightarrow_{\delta'}& \ ((U, p), (s_{init}, r), (L, q'), ((q, s), s')) \\
t^{(2)} &= ((L, q'), ((q, s), s'), (U, p), (s_{init}, r)) &\rightarrow_{\delta'}& \ ((L, q'), ((q, s), s'), (L, p'), (s_{ponr}, r')) \\
t^{(3)} &= ((L, p'), (s_{ponr}, r'), (L, q'), ((q, s), s')) &\rightarrow_{\delta'}& \ ((L, p'), (s_{ponr}, r'), (U, q'), (s_{init}, s')) \\
t^{(4)} &= ((x, y), (s_{init}, s'), (L, p'), (s_{ponr}, r')) &\rightarrow_{\delta'}& \ ((x, y), (s_{init}, s'), (U, p'), (s_{init}, r')) \\
& && \quad \text{for every } (x, y) \in Q' \\
t^{(5)} &= ((x, y), (v, w), (L, q'), ((q, s), s')) &\rightarrow_{\delta'}& \ ((x, y), (v, w), (U, q), (s_{init}, s)) \\
& && \quad \text{for every } (x, y) \in Q' \\
& && \quad \text{and } (v, w) \in S' \setminus \{(s_{ponr}, r')\}
\end{aligned}
$$

This simulation follows the same ideas as the Simulation 10. The only difference is the edge state of the original protocol that needs to be taken into account by transitions of the simulation and that needs to be backed up for possible future resets. We can now give a translation similar to Definition 14 required for our line of argumentation.

Definition 25 (Translation of Mediated Configurations)**.** Let $P = (Q, \Sigma, S, s_0, I, O, \delta)$ be a mediated population protocol and $P' = (Q', \Sigma', S', s_0', I', O', \delta')$ a immediate observation mediated population protocol constructed from P using Simulation 24. By $[\![\cdot]\!] : POP_\mathrm{M}(Q) \to POP_\mathrm{M}(Q')$ we denote the translation of configurations $C \in POP_\mathrm{M}(Q)$ in the population protocol into configurations $D \in POP_\mathrm{M}(Q')$ from the mediated population protocol. This translation is defined as follows:

$$
\left[\!\!\left[
\begin{pmatrix}
q_1 & s_{2,1} & \cdots & s_{n,1} \\
s_{1,2} & q_2 & \ddots & \vdots \\
\vdots & \ddots & \ddots & s_{n,n-1} \\
s_{1,n} & \cdots & s_{n-1,n} & q_n
\end{pmatrix}
\right]\!\!\right]
=
\begin{pmatrix}
(U, q_1) & (s_{init}, s_{2,1}) & \cdots & (s_{init}, s_{n,1}) \\
(s_{init}, s_{1,2}) & (U, q_2) & \ddots & \vdots \\
\vdots & \ddots & \ddots & (s_{init}, s_{n,n-1}) \\
(s_{init}, s_{1,n}) & \cdots & (s_{init}, s_{n-1,n}) & (U, q_n)
\end{pmatrix}
$$

With this simulation and translation we conjecture that each mediated population protocol can be simulated by a immediate observation mediated population protocol that shares several attributes, especially computing the same predicates.

Conjecture 26. For any mediated population protocol $P = (Q, \Sigma, S, s_0, I, O, \delta)$, the immediate observation mediated population protocol $P' = (Q', \Sigma', S', s_0', I', O', \delta')$ constructed from P using Simulation 24 calculates the same semilinear predicate.

5 Conclusion and future work

We have given a proof for the model of immediate observation mediated population protocols to compute all semilinear predicates. Thus they are as least as powerful in computation as population protocols. Additionally, we have given arguments why we believe this model is even equivalent to the model of mediated population protocols with two-way communication. Consequently allowing the initiator of a transition to change its agent and edge states does not contribute to the computational power. The proof of our Conjecture 26 can hopefully be done in our future research.

Our approach asks for a simulation and a translation which might seem overly complicated for a proof of equal computational power. But additionally several other attributes, besides the computation of the same predicate, carry over from the one protocol to the other if our Simulation 10 and translation from Definition 14 are used. Consider for example livelock freedom, i.e. no configuration is reached that has no successor besides itself. The simulation can reach a livelock iff the original protocol can reach such a configuration. This can easily be derived from the operational correspondence in Definition 15. In the context of protocols computing some predicate, a livelock is only possible if an output is reached. Otherwise the requirements for a well-specified protocol are not met. A livelock can be a desirable state as the computation can clearly be stopped in such a configuration. If the protocol does something else than computing a predicate, livelocks can be even more important to be reached or avoided, depending on the situation.

Another example is the analysis of a required communication structure. Whereas some protocols need a full interaction graph to carry out a computation, a path structure would suffice for others to get a correct result. The interaction graphs supporting a protocol do also support its simulation.

As a last example consider failure resistance [9]. If a protocol is designed to tolerate a certain number and type of faults, the simulation of this protocol could be capable of a comparable behaviour. This however depends on the type of failure and the chosen strategy to handle it. Crash failures, where an agent may leave the population at any time, should be manageable in the simulation with the same mechanisms as the original protocol did. Message losses could lead to new problems in the simulation

like deadlocked communication partners. Some error handling and error masking strategies could lead to the number of failures tolerable by the simulation being reduced in contrast to the original protocol.

A study on desirable attributes and how they carry over from one protocol to another by our simulation is something we wish to address in the future.

References

[1] Dan Alistarh, Rati Gelashvili & Milan Vojnović (2015): *Fast and Exact Majority in Population Protocols*. In: *Proceedings of the 2015 ACM Symposium on Principles of Distributed Computing*, PODC '15, ACM, New York, NY, USA, pp. 47–56, doi:10.1145/2767386.2767429.

[2] Dana Angluin, James Aspnes, Zoë Diamadi, Michael J. Fischer & René Peralta (2004): *Computation in networks of passively mobile finite-state sensors*. In: *Proceedings of the twenty-third annual ACM symposium on Principles of distributed computing*, ACM, pp. 290–299, doi:10.1145/1011767.1011810.

[3] Dana Angluin, James Aspnes, Zoë Diamadi, Michael J. Fischer & René Peralta (2006): *Computation in networks of passively mobile finite-state sensors*. *Distributed Computing* 18(4), pp. 235–253, doi:10.1007/s00446-005-0138-3.

[4] Dana Angluin, James Aspnes, David Eisenstat & Eric Ruppert (2007): *The computational power of population protocols*. *Distributed Computing* 20(4), pp. 279–304, doi:10.1007/s00446-007-0040-2.

[5] Dana Angluin, James Aspnes, Michael J. Fischer & Hong Jiang (2008): *Self-stabilizing Population Protocols*. *ACM Trans. Auton. Adapt. Syst.* 3(4), pp. 13:1–13:28, doi:10.1145/1452001.1452003.

[6] James Aspnes & Eric Ruppert (2009): *An Introduction to Population Protocols*, pp. 97–120. Springer Berlin Heidelberg, Berlin, Heidelberg, doi:10.1007/978-3-540-89707-1_5.

[7] Shantanu Das, Giuseppe Antonio Di Luna, Paola Flocchini, Nicola Santoro & Giovanni Viglietta (2017): *Mediated Population Protocols: Leader Election and Applications*. In T V Gopal, Jäger Gerhard & Silvia Steila, editors: *Theory and Applications of Models of Computation*, Springer International Publishing, Cham, pp. 172–186, doi:10.1007/978-3-319-55911-7_13.

[8] Shantanu Das, Giuseppe Antonio Di Luna, Paola Flocchini, Nicola Santoro & Giovanni Viglietta (2017): *Mediated Population Protocols: Leader Election and Applications*. In: *Theory and Applications of Models of Computation*, Lecture Notes in Computer Science, Springer, Cham, pp. 172–186, doi:10.1007/978-3-319-55911-7_13.

[9] Carole Delporte-Gallet, Hugues Fauconnier, Rachid Guerraoui & Eric Ruppert (2006): *When Birds Die: Making Population Protocols Fault-tolerant*. In: *Proceedings of the Second IEEE International Conference on Distributed Computing in Sensor Systems*, DCOSS'06, Springer-Verlag, Berlin, Heidelberg, pp. 51–66, doi:10.1007/11776178_4.

[10] Javier Esparza, Pierre Ganty, Jérôme Leroux & Rupak Majumdar (2017): *Verification of population protocols*. *Acta Informatica* 54(2), pp. 191–215, doi:10.1007/s00236-016-0272-3.

[11] Daniele Gorla (2010): *Towards a unified approach to encodability and separation results for process calculi*. *Information and Computation* 208(9), pp. 1031 – 1053, doi:10.1016/j.ic.2010.05.002.

[12] Othon Michail, Ioannis Chatzigiannakis & Paul G. Spirakis (2011): *Mediated population protocols*. *Theoretical Computer Science* 412(22), pp. 2434–2450, doi:10.1016/j.tcs.2011.02.003.